From Network Traces to System Responses: Opaquely Emulating Software Services

Miao Du*, Steve Versteeg†, Jean-Guy Schneider*, John Grundy*, Jun Han*
*School of Software and Electrical Engineering
Swinburne University of Technology, Hawthorn, Victoria 3122, Australia
Email: {miaodu,jschneider,jgrundy,jhan}@swin.edu.au
†CA Technologies, Melbourne, Victoria 3004, Australia
Email: steve.versteeg@ca.com

arXiv:1510.01421v2 [cs.SE] 19 Apr 2016

***Abstract*—Enterprise software systems make complex interactions with other services in their environment. Developing and testing for production-like conditions is therefore a challenging task. Prior approaches include emulations of the dependency services using either explicit modelling or record-and-replay approaches. Models require deep knowledge of the target services while record-and-replay is limited in accuracy. We present a new technique that improves the accuracy of record-and-replay approaches, without requiring prior knowledge of the services. The approach uses multiple sequence alignment to derive message prototypes from recorded system interactions and a scheme to match incoming request messages against message prototypes to generate response messages. We introduce a modified Needleman-Wunsch algorithm for distance calculation during message matching, wildcards in message prototypes for high variability sections, and entropy-based weightings in distance calculations for increased accuracy. Combined, our new approach has shown greater than 99% accuracy for four evaluated enterprise system messaging protocols.**

I. Introduction

Software systems are becoming ever increasing interconnected. Service emulation supports developing and testing a software system, independent of the other systems on which it depends [1]. In a typical deployment scenario, an enterprise system might interact with many other systems, such as a mainframe, directory servers, databases and other types of software services. In order to check whether the enterprise system will function correctly in terms of its interactions with these services, it is necessary to test it in as realistic an environment as possible, prior to the actual deployment. Getting access to the actual production environment for testing is not possible due to the risk of disruption. Large organisations often have a test environment, which is a close replication of their production environment, but this is very expensive. Furthermore the test environment is in high demand, so software developers will have only limited access to it. Enabling developers to have continuous access to production-like conditions to test their application is an important part of Development-Operations (known as 'DevOps') [2][3].

One popular approach that developers use to test their application's dependence on other systems is to install the other systems on virtual machines (such as VMware) [4]. However virtual machines are time consuming to configure and maintain. Furthermore the configuration of the systems running on the virtual machine is likely to be different to the production environment. An alternative that is gaining traction is service emulation, where models of services are emulated - sometimes into the many thousands of service instances - to provide more realistic scale and less complicated configuration [5]. However, existing approaches to service emulation rely on system experts explicitly modelling the target services and this requires very detailed knowledge of message protocols and structure. This is often infeasible if the required knowledge is unavailable for the wide range of services, especially legacy services, in a real deployment environment [6], and is besides very time consuming and error-prone.

Our aim is to develop an automated approach to service emulation which uses no explicit knowledge of the services, their message protocols and structures and yet can simulate – to a high degree of accuracy and scale – a realistic enterprise system deployment environment. To achieve this we need an automated, accurate, efficient and robust method for service emulation that derives service responses from collected actual network traces between an enterprise system and a service – the focus of this paper.

Our previous work [7] relied on searching the collected message traces to automatically generate service responses. Unfortunately this was limited to text based protocols and became inefficient for large message traces. The method we employed was sensitive to mismatching by payload versus operation type (i.e. was not robust.) Our subsequent work [8] used clustering to improve efficiency but decreased accuracy. The aim of the research described in this paper is to achieve both higher accuracy and efficiency through the use of descriptive and robust message prototypes which can be applied to both binary and text protocols. Our approach uses:

- a multiple sequence alignment algorithm to derive message prototypes,
- wildcards in message prototypes for the sections with high variability,
- a modified Needleman-Wunsch algorithm for message distance calculations, and
- entropy-based weightings in distance calculations for increased accuracy.

We present a set of experiments with four enterprise system

messaging protocols, including binary and textual LDAP, a binary Mainframe protocol, a textual JSON (Twitter) and SOAP services. These experiments show overall a greater than 99% accuracy in the generated response messages for the four protocols tested. Additionally they show efficient performance in generating the emulated service responses, enabling scaling within an emulated deployment environment.

II. Motivation and Related Work

Consider an example enterprise software system which is to be integrated with other systems in its production environment as part of its operation. The other systems (called services) include a legacy mainframe program, a directory server and a web service. The behaviour of the example enterprise software system depends on the responses it receives from these dependency services. The directory server uses a proprietary protocol which is poorly documented. It has operation types add, search, modify and delete, each with various payloads and response messages. The operation type is encoded in a single character. A small sample set of request and response messages (a *transaction library*) for the directory service is shown in Table I. Such message traces are captured by either network monitoring tools, such as Wireshark [9], or proxies. In the following sections we use this sample transaction library to illustrate how our method works. This transaction library is from a fictional protocol that has some similarities to the widely used LDAP protocol [10], but is simplified to make our running example easier to follow.

The enterprise system under test needs to interact with this service, to search for user identities, register new identities, deregister identities and so on. Other systems in the environment share the same service. For complex deployment environments we want to extensively test the enterprise system as to its protocol conformance, robustness and scaling.

The most common approach to providing a testing environment for enterprise systems using such a messaging protocol is to use virtual machines [11]. Implementations of the services are deployed on virtual machines and communicated with by the system under test. Major challenges with this approach include configuration complexity [12] and the need to maintain instances of each and every service type in multiple configurations [13]. Recently, cloud-based testing environments have emerged to mitigate some of these issues [14].

Emulated testing environments for enterprise systems, relying on service models, are another approach. When sent messages by the enterprise system under test, these respond with approximations of "real" service response messages [15]. Kaluta [16] is proposed to provision emulated testing environments. Challenges with these approaches include developing the models, lack of precision in the models, especially for complex protocols, and ensuring robustness of the models under diverse loading conditions [17]. To assist developing reusable service models, approaches either reverse engineer message structures [18][19][20], or discover [21][22] processes for building behavioral models [23]. While these allow engineers to develop more precise models, none of them can automate interaction between enterprise system under test and the emulated dependency services.

Recording and replaying message traces is an alternative approach. This involves recording request messages sent by the enterprise system under test to real services and response messages from these services, and then using these message traces to 'mimic' the real service response messages in the emulation environment [24]. Some approaches combine record-and-replay with reverse-engineered service models. CA Service Virtualization [25] is a commercial software tool, which can emulate the behaviour of services. The tool assumes knowledge of protocol message structures to model services and mimic interactions automatically.

III. Approach

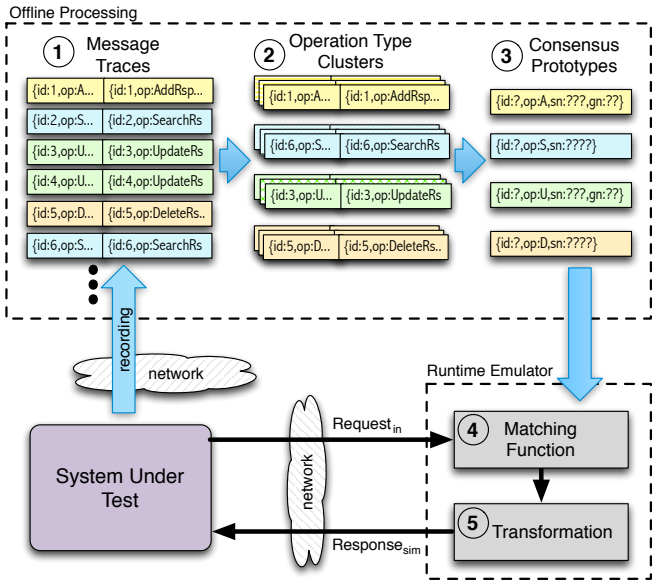

Fig. 1: System overview

Our goal is to produce an emulation environment for enterprise system testing that uses message trace recordings collected a priori to produce a response on behalf of a service when invoked by a system-under-test at runtime.

The general approach is to cluster the trace recordings into groups of similar messages and then formulate a single representation for the request messages of each cluster. This accelerates runtime performance by enabling incoming requests from the system under test to be compared only to the cluster representations, rather than the entire transaction library. Previous work [8] selected the cluster *centroid*[1] request as the representative. However accuracy decreased as the information from the other requests in the cluster was discarded. Our current aim is to improve accuracy, but preserve efficiency, by generating *prototypes* which capture the common features of the range of requests in each cluster.

[1] The centroid is the transaction with the minimised total distance from the other transactions in the cluster.

Index	Request	Response
1	{id:1,op:S,sn:Du}	{id:1,op:SearchRsp,result:Ok,gn:Miao,sn:Du,mobile:5362634}
13	{id:13,op:S,sn:Versteeg}	{id:13,op:SearchRsp,result:Ok,gn:Steve,sn:Versteeg,mobile:9374723}
24	{id:24,op:A,sn:Schneider,mobile:123456}	{id:24,op:AddRsp,result:Ok}
275	{id:275,op:S,sn:Han}	{id:275,op:SearchRsp,result:Ok,gn:Jun,sn:Han,mobile:33333333}
490	{id:490,op:S,sn:Grundy}	{id:490,op:SearchRsp,result:Ok,gn:John,sn:Grundy,mobile:44444444}
2273	{id:2273,op:S,sn:Schneider}	{id:2273,op:SearchRsp,result:Ok,sn:Schneider,mobile:123456}
2487	{id:2487,op:A,sn:Will}	{id:2487,op:AddRsp,result:Ok}
3106	{id:3106,op:A,sn:Hine,gn:Cam,Postcode:33589}	{id:3106,op:AddRsp,result:Ok}

TABLE I: Directory Service Interaction Library Example

The approach for generating and using our prototypes (depicted in figure 1) has the following steps:

1) We collect message traces of communications between a client and the real target service and store them in *a transaction library* with a transaction being a request-response pair.
2) We then cluster the transaction library, with the goal of grouping transactions by operation type. We do not consider the state of service under which the requests are issued. (This may give lower accuracy but is still useful in many cases as discussed in Section V.)
3) We then derive a *request consensus prototype* (a message pattern) for each operation type cluster by:
   a) Aligning all request messages of each cluster to reveal their common features for each operation type by inserting gap characters;
   b) Extracting a consensus character sequence from the aligned request messages (i.e., the request message prototype) by selecting or deleting a character at each character position;
   c) Calculating positional weightings to prioritise different sections of each request consensus prototype.
4) At runtime, for an incoming request message, we use *a matching distance calculation* technique to select the nearest matching request consensus prototype.
5) We then perform dynamic substitutions on a specifically chosen response message from the identified operation type cluster to generate a final response message to be sent back to the enterprise system under test.

Steps 1 to 3 are performed offline to prepare request consensus prototypes. Steps 4 and 5 are performed at runtime to generate live responses to the system under test messages.

A. Needleman-Wunsch

The Needleman-Wunsch algorithm [26] is used at different steps during the offline and runtime processing. It is a dynamic programming algorithm which finds the globally optimal alignment for two sequences of symbols in $O(MN)$ time, where M and N are the lengths of the sequences. Needleman-Wunsch uses a progressive scoring function S, which gives an incremental score for each pair of symbols in the alignment. A different score will be given depending on whether the symbols are identical, different or a gap has been inserted.

B. Collecting and Clustering Interaction Messages

The first step is to record real message exchanges between a client (such as a previous version of the system under test) and the service we aim to emulate. A transaction is a request-response pair of the communications, where the service may respond with zero or more response messages to a request message from the client. We record requests and responses at the network level (using a tool such as Wireshark), recording the bytes in TCP packet payloads. This makes no assumptions about the message format of the service application.

Having recorded our transaction library, the next step is to group the transactions by operation type, but again without assuming any knowledge of the message format. To achieve this we use a distance function-based clustering technique. In our previous work [8] we considered multiple cluster distance functions and found that the response similarity (as measured by edit distance, calculated using the Needleman-Wunsch algorithm [26]) was the most effective method for grouping transactions of the same operation type. We therefore use the same technique in this work, grouping transactions by the similarity of their response messages. The clustering algorithm used was VAT [27], one among many alternatives, as we found it to be effective in our previous work [8].

Applying our clustering technique to the example transaction library yields two clusters, as shown in tables II and III, corresponding to search operations and add operations, respectively.

C. Aligning Request Messages

To formulate a representative cluster prototype, we first align the request messages in a cluster to determine the common features of these request messages. Then we extract the common features while accommodating variations. In aligning the request messages of a cluster, we adopt the multiple sequence alignment (MSA) technique originated from Biology. MSA was first used to align three or more biological sequences to reveal their structural commonalities [28][2]. Using MSA for revealing commonalities of interaction messages offers a number of advantages over other techniques, such as n-grams [31]. Firstly, it is able to *effectively* handle messaging protocols with single-byte operation fields (e.g. in the binary LDAP protocol) as well as multi-byte operation fields. Secondly, it does not

[2]For this reason this technique has also been widely used to reverse-engineer protocol message structures [29], [30].

Index	Request	Response
1	{id:1,op:S,sn:Du}	{id:1,op:SearchRsp,result:Ok,gn:Miao,sn:Du,mobile:5362634}
13	{id:13,op:S,sn:Versteeg}	{id:13,op:SearchRsp,result:Ok,gn:Steve,sn:Versteeg,mobile:9374723}
275	{id:275,op:S,sn:Han}	{id:275,op:SearchRsp,result:Ok,gn:Jun,sn:Han,mobile:33333333}
490	{id:490,op:S,sn:Grundy}	{id:490,op:SearchRsp,result:Ok,gn:John,sn:Grundy,mobile:44444444}
2273	{id:2273,op:S,sn:Schneider}	{id:2273,op:SearchRsp,result:Ok,sn:Schneider,mobile:123456}

TABLE II: Cluster 1 (search operations)

Index	Request	Response
24	{id:24,op:A,sn:Schneider,mobile:123456}	{id:24,op:AddRsp,result:Ok}
2487	{id:2487,op:A,sn:Will}	{id:2487,op:AddRsp,result:Ok}
3106	{id:1106,op:A,sn:Hine,gn:Cam,postalCode:33589}	{Id:1106,Msg:AddRsp,result:Ok}

TABLE III: Cluster 2 (add operations)

require a predetermined matching sequence width (e.g., n in n-gram).

In particular, we adopt *ClustalW*[32], a widely used heuristic technique for MSA. It is memory efficient and is shown to produce high accuracy alignments in polynomial computation time for empirical datasets (in contrast to the original NP-complete MSA technique [33]). ClustalW is a progressive MSA algorithm, where pairwise sequence alignment results are iteratively integrated into the multiple sequence alignment result. An overview of the ClustalW algorithm is as follows:

1) All $N(N-1)/2$ pairs of sequences are aligned to calculate their similarity ratio by using the Needleman-Wunsch algorithm, where N is the number of sequences.
2) A $N \times N$ distance matrix is built for capturing distance calculation.
3) A *guide tree* is constructed from the distance matrix by applying a neighbour-joining clustering algorithm. (A guide tree is a tree data structure which organises the similarities between sequences.)
4) The guide tree is used to guide a progressive alignment of sequences from the leaves to the root of the tree.

Figure 2 shows the multiple sequence alignment results of applying the $ClustalW$ algorithm to the example clusters from tables II and III. The MSA results are known as *profiles*. Gaps which were inserted during the alignment process are denoted with the '⋆' symbol. Note that the common sequences for the requests in each cluster have now been aligned.

```
{id:⋆⋆1⋆,op:S,sn:⋆⋆⋆⋆⋆⋆⋆Du}
{id:⋆⋆13,op:S,sn:⋆Versteeg}
{id:2273,op:S,sn:Schneider}
{id:275⋆,op:S,sn:⋆Han⋆⋆⋆⋆⋆}
{id:490⋆,op:S,sn:Grundy⋆⋆⋆}
```

(a) Cluster 1 alignment

```
{id:24⋆⋆,op:A,sn:Schne⋆⋆⋆⋆⋆ider⋆⋆,mo⋆bil⋆⋆⋆e:123456}
{id:2487,op:A,sn:W⋆⋆⋆⋆⋆⋆⋆⋆⋆⋆⋆⋆⋆⋆⋆⋆⋆⋆il⋆⋆⋆l⋆⋆⋆⋆⋆⋆⋆}
{id:3106,op:A,sn:Hi⋆ne,gn:Cameron,postalCode:3⋆3589}
```

(b) Cluster 2 alignment

Fig. 2: MSA results of the requests in tables II and III

D. Formulating the Request Consensus Prototype

Having derived the MSA profile for the request messages of each cluster, the next step is to extract the common features from the MSA profile into a single character sequence, which we call the *request consensus prototype*, to facilitate efficient runtime comparison with an incoming request message from the system under test.

From all the aligned request messages in a cluster, we derive a byte (or character) occurrence count table. Figure 3 graphically depicts byte frequencies at each position for the example alignment in Figure 2. Each column represents a position in the alignment result. The frequencies of the different bytes which occur at each position are displayed as a stacked bar graph.

Based on the byte occurrence table, we formulate the *request consensus prototype* by extending the concept of a *consensus sequence* [34] commonly used in summarising a MSA profile. A consensus sequence can be viewed as a sequence of *consensus symbols*, where the consensus symbol c_i is the most commonly occurring byte at the position i. In our extension, a request consensus prototype, $\mathbf{p}$, is calculated by iterating each byte position of the MSA profile, to calculate a *prototype symbol*, p_i, at each position, according to equation 1.

$$p_i = \begin{cases} c_i & \text{if } q(c_i) \geq f \wedge c_i \neq \star \\ \perp & \text{if } q(c_i) \geq \frac{1}{2} \wedge c_i = \star \\ \mathbf{?} & \text{otherwise} \end{cases} \quad (1)$$

Where $q(c_i)$ denotes the relative frequency at position i of the consensus symbol c_i, f is the relative frequency threshold, '⋆' denotes a gap, '**?**' is the 'wildcard' symbol and '$\perp$' represents a truncation. After calculating the prototype symbol for each position, any truncation symbols are then deleted from the request consensus prototype.

Introducing wildcards and truncations into the message prototypes allows us to distinguish between gaps and where there is no consensus. If the relative frequency $q(c_i)$ is at or above the threshold f, then we insert the consensus symbol

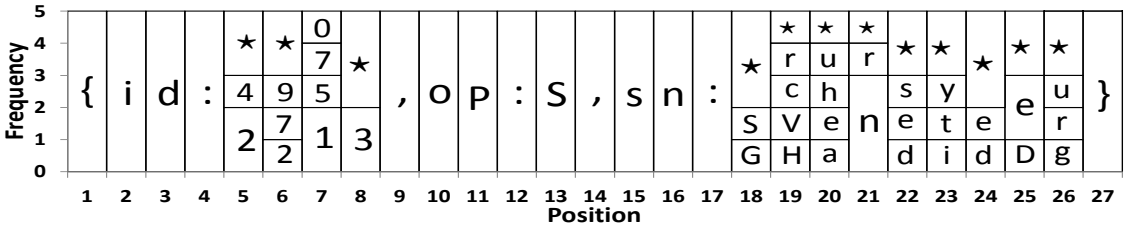

Fig. 3: Character Frequencies in the Alignment Result of Search Requests.

into our prototype (unless the consensus symbol is a gap.) If the relative frequency is below threshold, then we insert a wildcard. If the consensus symbol is a gap and it is in the majority, then we leave that position as empty (i.e. deleted). Wildcards allow us to encode where there are high variability sections of the message. In our experimentation, without truncations consensus sequences became artificially long as there tended to be many gaps. By truncating the gaps, the lengths of the prototypes become similar to the typical lengths of messages in the cluster. The consensus prototype for a cluster of request messages can differentiate stable positions from variant positions. Moreover, it can identify consensus symbols that can be utilised for matching.

Applying our request consensus prototype method, using a frequency threshold $f = 0.8$, to the example clusters in tables II and III yields the following results:

Consensus prototype for the search request cluster:

`{id:???,op:S,sn:???????}`

Consensus prototype for the add request cluster:

`{id:????,op:A,sn:????????????????l???????}`

(Note that the add prototype contains an 'l' from coincidentally aligning 'l's from 'mobile', 'Will' and 'postalCode'.)

E. Deriving Entropy-Based Positional Weightings

The final step of the offline analysis is to calculate weightings for each byte position in the consensus prototype to support the runtime distance matching process for request messages. In our service emulator, generating a response message of the correct operation type is more critical than the contents of the message payload. Thus we give a higher weighting to the sections of the message that likely relate to the operation type. To do this we make use of the observation that structure information (such as operation type) is more stable than payload information. We use entropy as a measurement of variability, and use it as the basis to calculate a weighting for each byte position of the consensus prototype.

Using the MSA profile for a cluster (from section III-C), we calculate the entropy for each column using the Shannon Index [35], as given by equation 2.

$$E_i = -\sum_{j=1}^{R} q_{ij} \log q_{ij} \tag{2}$$

Where E_i is the Shannon Index for the ith column, q_{ij} is the relative frequency of the jth character in the character set at column i and R is the total number of characters in the character set.

Since we wish to give a high weighting to stable parts of the message, for each column we invert the entropy by applying a scaling function of the form given in equation 3

$$w_i = \frac{1}{(1 + bE_i)^c} \tag{3}$$

Where w_i is the weighting for the ith column, E_i is the entropy of the ith column, and b and c are positive constants. The higher the values of b and c, the more higher entropy columns are deweighted. In our experiments we found the best results were obtained with $b = 1$ and $c = 10$. This allows structural information to strongly dominate in the matching process. Payload similarity is only used as a 'tie-breaker'.

Columns that correspond to gaps removed from the consensus prototype are also dropped from the weightings array. Table IV gives an example weightings array for the search request consensus prototype.

F. Entropy Weighted Runtime Matching for Request Messages

At runtime, the formulated request consensus prototypes are used to match incoming requests from the system under test. We extend the Needleman-Wunsch algorithm to calculate the matching distance between an incoming request and the request consensus prototype for each operation type (or cluster). We modify the Needleman-Wunsch scoring function, S (see equation 4), by giving a special score for alignments with wildcard characters and multiplying all scores by their corresponding positional entropy weights (from equation 3).

$$S(p_i, r_j) = \begin{cases} w_i M & \text{if } p_i = r_j \wedge p_i \neq \textbf{?} \\ w_i D & \text{if } p_i \neq r_j \wedge p_i \neq \textbf{?} \\ w_i X & \text{if } p_i = \textbf{?} \end{cases} \tag{4}$$

Where p_i is the ith character in the consensus prototype, r_j is the jth character in the incoming request, w_i is the weighting for the ith column, M and D are constants denoting the Needleman-Wunsch identical score and difference penalty, respectively, and X is the wildcard matching constant. In our experiments we used values $M = 1$, $D = -1$ and $X = 0$.

Using the modified scoring equation 4, we apply Needleman-Wunsch to align the consensus prototype with an incoming request, calculating an absolute alignment score, s.

The relative distance, denoted as d_{rel}, is calculated from the absolute alignment score to normalise for consensus prototypes of different lengths, different entropy weights and differ-

p	{	i	d	:	?	?	?	,	o	p	:	S	,	s	n	:	?	?	?	?	?	?	?	}
E	0	0	0	0	1.05	1.33	1.33	0	0	0	0	0	0	0	0	0	1.61	1.61	0.95	1.33	1.33	1.05	1.33	0
w	1	1	1	1	$\frac{1}{1342}$	$\frac{1}{4760}$	$\frac{1}{4760}$	1	1	1	1	1	1	1	1	1	$\frac{1}{14638}$	$\frac{1}{14638}$	$\frac{1}{796}$	$\frac{1}{4760}$	$\frac{1}{4760}$	$\frac{1}{1342}$	$\frac{1}{4760}$	1

TABLE IV: Example weightings for search consensus prototype calculated from the MSA profile in Figure 3 (using equation 3 constants $b = 1, c = 10$.)

ent numbers of wildcards. The relative distance is in the range 0 to 1, inclusive, where 0 signifies the best possible match with the consensus, and 1 represents the furthest possible distance. It is calculated according to equation 5:

$$d_{\text{rel}}(\mathbf{p}, \mathbf{r}) = 1 - \frac{s(\mathbf{p}, \mathbf{r}) - s_{\min}(\mathbf{p})}{s_{\max}(\mathbf{p})} \tag{5}$$

Where $s_{\max}$ is the maximum possible alignment score for the given consensus prototype and $s_{\min}$ is the minimum possible alignment score for the given consensus prototype. These are calculated in equations 6 and 7, respectively.

$$s_{\max}(\mathbf{p}) = \sum_{i=1}^{|\mathbf{p}|} S(\mathbf{p}_i, \mathbf{p}_i) \tag{6}$$

$$s_{\min}(\mathbf{p}) = \sum_{i=1}^{|\mathbf{p}|} S(\mathbf{p}_i, \varnothing) \tag{7}$$

Where $\varnothing$ is a special symbol, different to all of the characters in the consensus prototype.

The consensus prototype that gives the least distance to the incoming message is selected as the matching prototype, therefore identifying the matching transaction cluster.

As an example, suppose we receive an incoming add request with the byte sequence `{id:37,op:A,sn:Durand}`. Aligning the request against the search consensus prototype yields:

request: `{id:3*7,op:A,sn:*Durand}`
prototype: `{id:???,op:S,sn:???????}`

Using equation 5, the weighted relative distance is calculated to be 0.0715. In comparison, the add consensus prototype produces the alignment below with a relative distance of 0.068.

request: `{id:37**,op:A,sn:********D*u*r*an*d***}`
prototype: `{id:????,op:A,sn:????????????l???????}`

Consequently, the add prototype is the closest matching, causing the add cluster to be selected.

G. Response Generation through Dynamic Substitution

The final step in our approach is to send a customised response for the incoming request by performing some dynamic substitutions on a response message from the matched cluster. Here we use the symmetric field technique described in [7], where character sequences which occur in both the request and response messages of the chosen transaction are substituted with the corresponding characters from the live request in the generated response. We take the response from the centroid transaction [8] of the selected cluster and apply the symmetric field substitution.

In our example, the centroid transaction from the add cluster is given below. There is one symmetric field (boxed).

request: `{id:24,op:A,sn:Schneider,mobile:123456}`
response: `{id:24,op:AddRsp,result:Ok}`

After performing the symmetric field substitution the final generated response is: `{id:37,op:AddRsp,result:Ok}`

H. Implementation

We implemented our proposed system using Java, following the steps in the previous sections, using the architecture of Figure 1. Due to commercial agreements we are unable to release the source code of our implementation.

IV. Evaluation

We wanted to assess three key characteristics of our approach: *accuracy*, *efficiency* and *robustness*, aiming to answer the following research questions:

1) **RQ1** (**Accuracy**): Having *request consensus prototypes* formulated at the pre-processing stage, is our approach able to generate accurate, protocol-conformant responses? At runtime, can *positional weighting based matching* further improve the accuracy of our approach for generating such responses?
2) **RQ2** (**Efficiency**): Is our technique efficient enough to generate timely responses, even for large transaction libraries?
3) **RQ3** (**Robustness**): Given non-homogeneous clusters, which contain some messages of operation types different to the majority of messages, is our technique robust enough to generate accurate responses?

In order to answer these questions, we applied our technique on six message trace datasets of four case study protocols. We ran the technique on these datasets and assessed its accuracy, efficiency and robustness for each. We then compared these results to our previous techniques.

A. Case Study Protocols and Traces

Protocol	Binary/Text	Fields	#Ops.	#Transactions
IMS	binary	fixed length	5	800
LDAP	binary	length-encoded	10	2177
LDAP text (1)	text	delimited	10	2177
LDAP text (2)	text	delimited	6	1000
SOAP	text	delimited	6	1000
Twitter (REST)	text	delimited	6	1825

TABLE V: Experiment message trace datasets (available at [36])

We applied our techniques to four real-world protocols, *IMS* [37] (a binary mainframe protocol), *LDAP* [10] (a binary directory service protocol), *SOAP* [38] (a textual protocol, with an Enterprise Resource Planning (ERP) system messaging

system services), and *RESTful Twitter* [39] (a JSON protocol for the Twitter social media service).

We have used one message trace dataset for each of these protocols. In addition, *LDAP* has two extra datasets: a dataset with textual representation converted from the binary dataset (denoted by *LDAP text (1)*), and another textual dataset that was used in our prior work [8] (*LDAP text (2)*).

We chose these protocols because: (i) they are widely used in enterprise environments; (ii) some of them (LDAP text (2) and SOAP) were used in the evaluation of our prior work [7], [8]; (iii) they represent a good mix of text-based protocols (SOAP and RESTful Twitter) and binary protocols (IMS and LDAP); (iv) they use either fixed length, length encoding or delimiters to structure protocol messages;[3] and (v) each of them includes a diverse number of operation types, as indicated by the *Ops* column. The number of request-response interactions for each test case is shown as column *#Transactions* in Table V.

Our message trace datasets are available for download from our website [36].

B. 10-fold Cross-Validation Approach and Evaluation Criteria

Cross-validation [40] is a popular model validation method for assessing how accurately a predictive model will perform in practice. For the purpose of our evaluation, we applied the commonly used 10-fold cross-validation approach [41] to all the six case study datasets.

We randomly partitioned each of the original interaction datasets into 10 groups. Of these 10 groups, one group is considered to be the *evaluation group* for testing our approach, and the remaining 9 groups constitute the *training set*. This process is then repeated 10 times (the same as the number of groups), so that each of the 10 groups will be used as the evaluation group once.

When running each experiment with each trace dataset, we applied our approach to each request message in the *evaluation group*, referred to as the *incoming request*, to generate an *emulated response*. We recorded the time that our approach took to generate the response for each incoming request. This was used to evaluate the runtime efficiency of our consensus+weighting approach compared to other approaches.

Having generated a response for each incoming request, we utilised five criteria to determine its accuracy, thereby evaluating the ability of our approach to generate protocol-conformant responses. These five criteria are explained as follows using an example shown in Table VI. Consider the *incoming request* {id:37,op:A,sn:Durand} with the associated response {id:37,op:AddRsp,result:Ok} in the transaction library. The emulated response is considered to be:

1) **identical** if its character sequence is identical to the recorded (or expected) response (cf. Example (i) in Table VI);

[3]Given a protocol message, length fields or delimiters are used to convert its structure into a sequence of bytes that can be transmitted over the wire. Specifically, a length field is a number of bytes that show the length of another field, while a delimiter is a byte (or byte sequence) with a known value that indicates the end of a field.

(i)	Expected	{id:37,op:AddRsp,result:Ok}
	Generated	{id:37,op:AddRsp,result:Ok}
(ii)	Expected	{id:***37***,op:AddRsp,result:*Ok*}
	Generated	{id:***37***,op:AddRsp,result:***AlreadyExists***}
(iii)	Expected	{id:37,op:AddRsp,result:Ok}
	Generated	{id:***15***,op:AddRsp,result:***AlreadyExists***}
(iv)	Expected	{id:37,op:*AddRsp*,result:Ok}
	Generated	{id:15,op:***SearchRsp***,result:Ok,gn:Miao,sn:Du}
(v)	Expected	{id:37,op:*AddRsp*,result:Ok}
	Generated	{id:15,op:***AearchRsp***,result:Ok,gn:Miao,sn:Du}

TABLE VI: Examples for accuracy criteria: (i) Identical, (ii) Consistent, (iii) Protocol conformant, (iv) Well-formed, (v) Malformed.

2) **consistent** if it is of the expected operation type and has the critical fields in the payload replicated (cf. Example (ii) in Table VI where `id` is identical, but some of the other payload differs);
3) **protocol conformant** if its operation type corresponds to the expected response, but it differs in some payload information (cf. Example (iii) in Table VI where both the `id` and `result` tags differ);
4) **well-formed** if it is structured correctly (that is, it corresponds to one of the valid response messages), but has the wrong operation type (cf. Example (iv) in Table VI where the generated response is of a valid structure, but its operation type `op:SearchRsp` does not match the expected operation type `op:AddRsp`); and
5) **malformed** if it does not meet any of the above criteria (cf. Example (v) in Table VI where operation type `op:AearchRsp` in the generated response is invalid).

We further consider a generated response to be *valid* if it meets one of the first three criteria, that is, *identical*, *consistent* or *protocol conformant*. Otherwise, a generated response is considered to be *invalid*.

C. Compared Techniques

To our knowledge there is no other work that generates application layer responses directly from trace messages. Our comparison is therefore with our prior work, including the *Whole Library* [7] and the *Cluster Centroid* [8] approaches. Given an incoming request, the *Whole Library* approach searches the entire transaction library for its closest matching request to synthesize its response(s). This approach is effective in producing accurate responses. Experimental results revealed that more than 90% of generated responses were valid, as defined by the criteria in Section IV-B. However, it is generally too slow for real-time use.

The *Cluster Centroid* approach reduces the number of searches to the number of transaction library clusters. Therefore, it can generate responses for real time use, but with less accuracy. We use these two approaches as baselines to evaluate our new technique.

D. Evaluation Results

*1) **Accuracy (RQ1)**:* The accuracy evaluation is conducted to assess the capability of our approach for generating accurate

Protocol	Method		Accuracy Ratio	No.	Valid			Invalid	
					Identical	Consistent	Conformant	Well-formed	Malformed
IMS (binary)	Whole Library		75.25%	800	400	202	0	198	0
	Cluster Centroid		97.88%		400	383	0	17	0
	Consensus Only	f=0.5	***100%***		400	400	0	0	0
		f=0.8	***100%***		400	400	0	0	0
		f=1	***100%***		400	400	0	0	0
	Consensus + Weighting	f=0.5	***100%***		400	400	0	0	0
		f=0.8	***100%***		400	400	0	0	0
		f=1	***100%***		400	400	0	0	0
LDAP (binary)	Whole Library		94.12%	2177	248	17	1784	36	92
	Cluster Centroid		91.59%		263	17	1714	183	0
	Consensus Only	f=0.5	*87.74%*		268	14	1628	267	0
		f=0.8	*84.66%*		264	14	1565	334	0
		f=1	*79.97%*		259	14	1468	436	0
	Consensus + Weighting	f=0.5	*98.71%*		278	18	1853	28	0
		f=0.8	***99.95%***		278	18	1880	1	0
		f=1	*86.90%*		267	16	1609	285	0
LDAP text (1) (text)	Whole Library		100%	2177	1648	415	114	0	0
	Cluster Centroid		100%		811	1325	41	0	0
	Consensus Only	f=0.5	***100%***		1555	622	0	0	0
		f=0.8	***100%***		1555	622	0	0	0
		f=1	***100%***		1527	650	0	0	0
	Consensus + Weighting	f=0.5	***100%***		1559	618	0	0	0
		f=0.8	***100%***		1559	618	0	0	0
		f=1	***100%***		1559	618	0	0	0
LDAP text (2) (text)	Whole Library		92.9%	1000	927	2	0	71	0
	Cluster Centroid		73.4%		509	225	0	252	14
	Consensus Only	f=0.5	***100%***		808	192	0	0	0
		f=0.8	***100%***		808	192	0	0	0
		f=1	***100%***		808	192	0	0	0
	Consensus + Weighting	f=0.5	***100%***		808	192	0	0	0
		f=0.8	***100%***		808	192	0	0	0
		f=1	***100%***		808	192	0	0	0
SOAP (text)	Whole Library		100%	1000	77	923	0	0	0
	Cluster Centroid		100%		98	902	0	0	0
	Consensus Only	f=0.5	***100%***		96	904	0	0	0
		f=0.8	***100%***		96	904	0	0	0
		f=1	***100%***		96	904	0	0	0
	Consensus + Weighting	f=0.5	***100%***		96	904	0	0	0
		f=0.8	***100%***		96	904	0	0	0
		f=1	***100%***		96	904	0	0	0
Twitter (REST) (text)	Whole Library		**99.56%**	1825	150	994	673	7	1
	Cluster Centroid		99.34%		0	896	917	11	1
	Consensus Only	f=0.5	*99.34%*		0	893	920	11	1
		f=0.8	*99.34%*		0	893	920	11	1
		f=1	*99.34%*		0	893	920	11	1
	Consensus+ Weighting	f=0.5	*99.34%*		0	893	920	11	1
		f=0.8	*99.34%*		0	893	920	11	1
		f=1	*99.34%*		0	893	920	11	1

TABLE VII: Evaluation Results of Applying Consensus Sequence Prototype

responses. As discussed in Section III, our approach has two important features aimed at enhancing response accuracy, the *request consensus prototype* combined with *entropy-weighted distance calculation* at runtime. To measure the impact of these techniques we ran two separate sets of experiments, which are referred to as *Consensus Only* and *Consensus+Weighting*, respectively. In addition, for both sets of experiments we tested for the best pre-defined frequency threshold f, trying three different values.[4]

Table VII summarises the evaluation results of *Consensus Only*, *Consensus+Weighting*, and our prior work (i.e. *Whole Library* and *Cluster Centroid*) experiments for the six test datasets. The **Accuracy Ratio** column is calculated by dividing the number of valid generated responses by the total number of interactions tested. The last five columns give a more detailed breakdown of the different categories of valid and invalid responses generated.

Table VII shows that the combined *Consensus+Weighting* approach achieves the highest accuracy overall for the datasets tested. The combined approach achieves 100% accuracy for four of the datasets, and 99.95% and 99.34% for the remaining two (LDAP binary and Twitter, respectively). Twitter is the only case where the *Whole Library* approach is marginally better.

With respect to the impact for the frequency threshold f, the results show that allowing some tolerance (i.e. $f < 1$) in the

[4]As illustrated in Equation 1, a pre-defined frequency threshold is required to calculate the consensus sequence prototype. In bioinformatics, a number of investigations have been done for identifying the best threshold [34]. In our experiments, we selected the most popular 3 values, that is, 0.5, 0.8 and 1.

	No.	Whole Library	Cluster Centroid	Consensus +Weighting	Real System
IMS	800	470.99	4.94	3.24	518
LDAP	2177	835.91	2.77	2.88	28
LDAP text(1)	2177	1434.70	5.69	7.30	28
LDAP text(2)	1000	266.30	2.44	1.63	28
SOAP	1000	380.24	2.97	3.35	65
Twitter	1825	464.09	32.86	36.62	417

(a) Average Total Response Generation Time (ms)

	No.	Whole Library		Cluster Centroid		Consensus+W	
		M	S	M	S	M	S
IMS	800	460.78	10.2	3.67	1.27	2.68	0.56
LDAP	2177	828.95	6.96	2.38	0.39	2.60	0.28
LDAP text(1)	2177	1425.23	9.47	4.38	1.31	5.67	1.63
LDAP text(2)	1000	257.92	8.38	1.35	1.09	1.14	0.49
SOAP	1000	372.58	7.66	1.92	1.05	2.45	0.9
Twitter	1825	412.98	51.11	1.47	31.39	1.67	34.95

(b) Average Matching Time and Average Substitution Time. M represents the matching time, and S represents the substitution time. (ms)

TABLE VIII: Approach Efficiency Evaluation Results

multiple sequence alignment can yield better results. For the LDAP (binary) dataset the thresholds of $f = 0.5$ and $f = 0.8$ produced significantly higher accuracy than $f = 1$. For the other datasets the threshold had no impact on the accuracy. The general conclusion appears that the results are not very sensitive to the value of the threshold for most scenarios.

An interesting result is that the *Consensus+Weighting* approach has a higher accuracy than the *Whole Library* approach, even though the latter uses all the available data points from the trace library (for three datasets *Consensus+Weighting* is significantly more accurate, for two it has the same accuracy, for one it is slightly lower). The reason for the higher accuracy is that the *Consensus+Weighting* abstracts away the message payload information sections (using wildcards), so is less susceptible to matching a request to the wrong operation type but with the right payload information, whereas the *Whole Library* approach is susceptible to this type of error (note the well-formed but invalid responses for the *Whole Library* approach in Table VII).

The impact of the entropy weightings can only be observed for the LDAP binary dataset. For this test, the weightings significantly improve the accuracy results. For the other datasets, no impact from the weightings can be observed, as the consensus sequence prototype by itself (*Consensus Only*) already produces 99-100% accuracy.

*2) **Efficiency (RQ2)**:* Table VIIIa compares the average response generation time of the *Consensus+Weighting* approach with the *Whole Library* and *Cluster Centroid* approaches. Tests were run on an Intel Xeon E5440 2.83GHz CPUs with 24GB of main memory available. The times represent the average response generation time across all of the requests in the datasets. Table VIIIa also lists the average response times of the real services from which the original traces were recorded. In order to get a better insight of the runtime performance of our approach, we separately measured matching time and substitution time, results of which are presented in Table VIIIb.

The results show that the *Consensus+Weighting* approach is very efficient at generating responses, indeed much faster than the real services being emulated. The response generation time is comparable to the *Cluster Centroid* approach, being faster for some datasets, slower for others. Both of these approaches are about two orders of magnitude faster than the *Whole Library* approach. However, whereas the *Cluster Centroid* approach trades off accuracy for speed, the *Consensus+Weighting* has both high accuracy and speed.

Comparing the matching time versus the substitution time (shown in Table VIIIb) we can observe that the *Whole Library* approach consumes most of its time during the matching process (because a Needleman-Wunsch alignment is made with every request in the transaction library). *Consensus+Weighting* and *Cluster Centroid* have greatly reduced matching times. Twitter has unusually long substitution times, such that for the fast approaches, most time is spent performing the substitution. This is due to the Twitter responses being very long, causing the symmetric field identification (common substring search) to become time consuming.

The *Consensus+Weighting* generates responses faster than the real services being emulated. This is crucial for supporting testing of an enterprise system under test under realistic performance conditions (delays can be added to slow down the emulated response, but not the other way around). A major limitation of the *Whole Library* approach is that it cannot generate responses in a time which matches the real services for fast services (such as LDAP).

*3) **Robustness (RQ3)**:* Our final test is to evaluate whether our *Consensus+Prototype* approach is robust in generating accurate responses when the clustering process (from Section III-B) is imperfect. For this test we deliberately inject noise into our clusters, i.e. to create clusters where a fraction of the interaction messages are of different operation types. The noise ratios tested were 5%, 10% and 20%. We repeated the experiments with different frequency thresholds (i.e. 0.5, 0.8, and 1).

Table IX summarises the experimental results. The results show that having a frequency threshold below 1 has a very big impact on preserving the accuracy when the clustering is noisy. A threshold of $f = 0.5$ gives the best accuracy. When using this threshold, the accuracy stays above 97% for all datasets, when the noise ratio is 5%. As the noise ratio increases to 20%, the accuracy decreases significantly for binary LDAP, but stays high for the other datasets.

Overall this is a good result. For our approach to work best the clusters produced should be relatively clean, but there is tolerance for a small amount of noise. A noise ratio of 20% is considered very high. Our actual clustering process produced perfect separation (i.e. 0% noise) of interaction messages by operation type for the six datasets tested.

	5% Noise			10% Noise			20% Noise		
	Consensus+Weighting			Consensus+Weighting			Consensus+Weighting		
	f = 0.5	f = 0.8	f = 1	f = 0.5	f = 0.8	f = 1	f = 0.5	f = 0.8	f = 1
IMS	*99.38%*	*99.75%*	*72.25%*	*95.75%*	*98.63%*	*68.5%*	*99.75%*	*96.5%*	*68.13%*
LDAP	*97.01%*	*97.29%*	*33.49%*	*87.09%*	*79.83%*	*49.93%*	*73.81%*	*63.67%*	*42.03%*
LDAP text (1)	*100%*	*100%*	*83.74%*	*100%*	*100%*	*84.20%*	*100%*	*100%*	*86.50%*
LDAP text (2)	*100%*	*100%*	*100%*	*100%*	*100%*	*100%*	*100%*	*100%*	*100%*
SOAP	*98.2%*	*98.2%*	*91.5%*	*100%*	*100%*	*98.2%*	*100%*	*100%*	*97.7%*
Twitter (REST)	*99.34%*	*99.34%*	*22.41%*	*97.59%*	*95.01%*	*32.60%*	*96.71%*	*94.13%*	*29.00%*

TABLE IX: Response Accuracy for Clusters with Noisy Data

E. Industry Validation

The proposed service emulation technique has been integrated into CA Technologies' commercial product: CA Service Virtualization. An earlier version of the technique [42] was released as a new feature, named *opaque data processing*, in version 8.0 of the product and has been sold to customers [43]. Opaque data processing has been used at customer sites to successfully emulate services of protocols not otherwise supported by the product. The emulated protocols included a proprietary extension of IMS and Sun's ONC/RPC protocol. The present technique is on the backlog for future versions.

F. Threats to Validity

We have identified some threats to validity which should be taken into consideration when generalising our experimental results:

- Our evaluation was performed on six datasets from four protocols. Given the great diversity in message protocols, further testing should be performed on other message protocols.
- The datasets were obtained by randomly generating client requests for services of different protocols. Some real system interactions are likely to be more complicated than those of our datasets. Further testing on real system interactions are warranted.

V. Discussion and Future Work

We have developed an approach for automatically generating service responses from message traces which requires no prior knowledge of message structure or decoders or message schemas. Our approach of using multiple sequence alignment to automatically generate consensus prototypes for the purpose of matching request messages is shown to be accurate, efficient and robust. Wildcards in the prototypes allow the stable and unstable parts of the request messages for the various operation types to be separated. Rather than using the prototypes directly for strict matching (such as using it as a regular expression) we instead calculate matching distance through a modified Needleman-Wunsch alignment algorithm. Since we look for the closest matching prototype, the method is robust even if the prototypes are imperfect. Moreover, this process can match requests which are slightly different to the prototypes, or are of different length to the prototypes. This allows the system to handle requests which are outside of the cases directly observed in the trace recordings. Weighting sections of the prototype with different importance based on the entropy, further improves the matching accuracy.

Our experimental results using the 6 message trace datasets demonstrate that our approach is able to automatically generate accurate responses in real time for most cases. Moreover, we can also see that our approach can also generate accurate responses from imperfect message clusters that contain a small number of messages of different operation types.

One limitation in our current approach is the lack of diversity in the responses generated. We are working on an approach to identify common patterns of all responses in a cluster. A possible solution is to apply multiple sequence alignment to response messages, to distinguish stable positions from variable positions. The variable parts of responses could then be stochastically generated. To further improve the robustness of our approach to noisy clustering, we will utilise outlier detection techniques [44] to remove outliers of clusters before applying the alignment method.

Our approach does not consider the service state history in formulating responses. In practice a stateless model is sufficient in many cases. For example: (i) when the emulation target service is stateless, or (ii) when the testing scenario does not contain two equivalent requests, requiring different state affected responses, or (iii) where the testing scenario does not require highly accurate responses (e.g. performance testing.) To address this limitation, an avenue of future exploration is to process mine the operation sequences to discover stateful models.

VI. Conclusion

We have developed a new technique for automatically generating realistic response messages from network traces for enterprise system emulation environments that outperforms current approaches. We use the bioinformatics-inspired multiple sequence alignment algorithm to derive message prototypes, adding wildcards for high variability sections of messages. A modified Needleman-Wunsch algorithm is used to calculate message distance and entropy weightings are used in distance calculations for increased accuracy. Our technique is able to automatically separate the payload and structural information in complex enterprise system messages, making it highly robust. We have shown in a set of experiments with four enterprise system messaging protocols a greater than 99% accuracy for the four protocols tested. Additionally they show efficient emulated service response performance enabling scaling within an emulated deployment environment.

REFERENCES

[1] M. J. Rutherford, A. Carzaniga, and A. L. Wolf, "Evaluating test suites and adequacy criteria using simulation-based models of distributed systems," *IEEE Transactions on Software Engineering*, vol. 34, no. 4, pp. 452–470, 2008.

[2] F. Erich, C. Amrit, and M. Daneva, "Report: Devops literature review," University of Twente, Tech. Rep., 2014. [Online]. Available: http://www.utwente.nl/bms/iebis/staff/amrit/devopsreport.pdf

[3] L. Bass, I. Weber, and L. Zhu, *DevOps: A Software Architect's Perspective*. Addison-Wesley Professional, 2015.

[4] J. Sugerman, G. Venkitachalam, and B.-H. Lim, "Virtualizing i/o devices on vmware workstation's hosted virtual machine monitor," in *Proceedings of the General Track: 2001 USENIX Annual Technical Conference (USENIX 2001)*, Boston, Massachusetts, USA, 2001, pp. 1–14.

[5] C. Hine, J.-G. Schneider, J. Han, and S. Versteeg, "Scalable emulation of enterprise systems," in *Proceedings of the 20th Australian Software Engineering Conference (ASWEC 2009)*, Gold Coast, Australia, 2009, pp. 142–151.

[6] S. Ghosh and A. P. Mathur, "Issues in Testing Distributed Component-Based Systems," in *Proceedings of the 1st International ICSE Workshop on Testing Distributed Component-Based Systems (ICSE 1999)*, Los Angeles, California, USA, 1999.

[7] M. Du, J.-G. Schneider, C. Hine, J. Grundy, and S. Versteeg, "Generating service models by trace subsequence substitution," in *Proceedings of the 9th International ACM Sigsoft Conference on Quality of Software Architectures (QoSA 2013)*, Vancouver, British Columbia, Canada, 2013, pp. 123–132.

[8] M. Du, S. Versteeg, J.-G. Schneider, J. C. Han, and J. Grundy, "Interaction traces mining for efficient system responses generation," in *Proceedings of the 2nd International Workshop on Software Mining (SoftMine 2013)*, Palo Alto, CA, USA, 2013, pp. 1–8.

[9] G. Combs, "Wireshark," http://wireshark.org, 1998–2015.

[10] J. Sermersheim, "Lightweight Directory Access Protocol (LDAP): The Protocol," 6 2006, RFC 4511.

[11] P. Li, "Selecting and Using Virtualization Solutions: our Experiences with VMware and VirtualBox," *Journal of Computing Sciences in Colleges*, vol. 25, no. 3, pp. 11–17, 2010.

[12] P. Godefroid, "Micro execution," in *in the proceedings of the 36th International Conference on Software Engineering (ICSE 2014)*, Hyderabad, India, 2014, pp. 539–549.

[13] J. Grundy, Y. Cai, and A. Liu, "Softarch/mte: Generating distributed system test-beds from high-level software architecture descriptions," *Automated Software Engineering*, vol. 12, no. 1, pp. 5–39, Jan. 2005. [Online]. Available: http://dx.doi.org/10.1023/B:AUSE.0000049207.62380.74

[14] T. Banzai, H. Koizumi, R. Kanbayashi, T. Imada, T. Hanawa, and M. Sato, "D-cloud: Design of a software testing environment for reliable distributed systems using cloud computing technology," in *Proceedings of the 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing (CCGrid 2010)*, Melbourne, Australia, 2010, pp. 631–636.

[15] A. Bertolino, G. De Angelis, L. Frantzen, and A. Polini, "Model-based generation of testbeds for web services," in *Testing of Software and Communicating Systems*, 2008, vol. 5047, pp. 266–282.

[16] C. Hine, "Emulating enterprise software environments," PhD thesis, Swinburne University of Technology, Faculty of Information and Communication Technologies, 2012.

[17] J. Sun and T. Mannisto, "Usefulness evaluation of simulation in server system testing," in *Proceedings of the 36th IEEE Computer Software and Applications Conference (COMPSAC 2012)*, Izmir, Turkey, 2012, pp. 158–163.

[18] G. Wondracek, P. M. Comparetti, C. Kruegel, E. Kirda, and S. S. S. Anna, "Automatic network protocol analysis," in *Proceedings of the Network and Distributed System Security Symposium (NDSS 2008)*, California, USA, 2008, pp. 1–14.

[19] W. Cui, J. Kannan, and H. J. Wang, "Discoverer: Automatic protocol reverse engineering from network traces," in *Proceedings of the 16th USENIX Security Symposium (Security 2007)*, Boston, MA, USA, 2007, pp. 199–212.

[20] Y. Wang, X. Yun, M. Z. Shafiq, L. Wang, A. X. Liu, Z. Zhang, D. Yao, Y. Zhang, and L. Guo, "A semantics aware approach to automated reverse engineering unknown protocols," in *Network Protocols (ICNP), 2012 20th IEEE International Conference on*. IEEE, 2012, pp. 1–10.

[21] J. De Weerdt, J. Vanthienen, B. Baesens *et al.*, "Active trace clustering for improved process discovery," *IEEE Transactions on Knowledge and Data Engineering*, vol. 25, no. 12, pp. 2708–2720, 2013.

[22] R. Jagadeesh Chandra Bose and W. M. van der Aalst, "Process diagnostics using trace alignment: opportunities, issues, and challenges," *Information Systems*, vol. 37, no. 2, pp. 117–141, 2012.

[23] D. Lo, L. Mariani, and M. Pezzè, "Automatic steering of behavioral model inference," in *Proceedings of the 7th joint meeting of the European Software Engineering Conference and the ACM SIGSOFT International Symposium on Foundations of Software Engineering (ESEC/FSE 2009)*, Amsterdam, The Netherlands, 2009, pp. 345–354.

[24] W. Cui, V. Paxson, N. C. Weaver, and R. H. Katz, "Protocol-independent adaptive replay of application dialog," in *Proceedings of the 13th Annual Network and Distributed System Security Symposium (NDSS 2006)*, San Diego, California, USA, 2006, pp. 1–15.

[25] J. Michelsen, "Key Capabilities of a Service Virtualization Solution," October 2012, White Paper, CA Technologies. [Online]. Available: http://www.ca.com/us/~/media/files/whitepapers/ca-svcapabilities-wp-oct2012.aspx

[26] S. B. Needleman and C. D. Wunsch, "A general method applicable to the search for similarities in the amino acid sequence of two proteins," *Journal of Molecular Biology*, vol. 48, no. 3, pp. 443–453, 1970.

[27] J. C. Bezdek and R. J. Hathaway, "Vat: A tool for visual assessment of (cluster) tendency," in *Proceedings of the 2002 International Joint Conference on Neural Networks (IJCNN 2002)*, Honolulu, Hawaii, USA, 2002, pp. 2225–2230.

[28] R. Durbin, *Biological sequence analysis: probabilistic models of proteins and nucleic acids*. Cambridge university press, 1998.

[29] P. M. Comparetti, G. Wondracek, C. Kruegel, and E. Kirda, "Prospex: Protocol Specification Extraction," in *Proceedings of the 30th IEEE Symposium on Security and Privacy (S&P 2009)*, Oakland, CA, USA, 2009, pp. 110–125.

[30] M. Beddoe, "The protocol informatics project," *Toorcon*, vol. 4, p. 4, 2004.

[31] Y. Wang, X. Yun, M. Z. Shafiq, L. Wang, A. X. Liu, Z. Zhang, D. Yao, Y. Zhang, and L. Guo, "A semantics aware approach to automated reverse engineering unknown protocols," in *Proceedings of the 20th IEEE International Conference on Network Protocols (ICNP 2012)*, Austin, TX, USA, 2012, pp. 1–10.

[32] J. D. Thompson, D. G. Higgins, and T. J. Gibson, "CLUSTAL W: improving the sensitivity of progressive multiple sequence alignment through sequence weighting, position-specific gap penalties and weight matrix choice," *Nucleic Acids Research*, vol. 22, no. 22, pp. 4673–4680, 1994.

[33] L. Wang and T. Jiang, "On the complexity of multiple sequence alignment," *Journal of computational biology*, vol. 1, no. 4, pp. 337–348, 1994.

[34] W. H. Day and F. McMorris, "Critical comparison of consensus methods for molecular sequences," *Nucleic Acids Research*, vol. 20, no. 5, pp. 1093–1099, 1992.

[35] C. E. Shannon, "A mathematical theory of communication," *The Bell System Technical Journal*, pp. 379–423,623–656, 1948.

[36] M. Du and S. Versteeg, "Interactive Message Traces from Application Layer Protocols," 9 2014. [Online]. Available: http://quoll.ict.swin.edu.au/doc/message_traces.html

[37] R. Long, M. Harrington, R. Hain, and G. Nicholls, *IMS Primer*. IBM International Technical Support Organisation, 2000. [Online]. Available: http://www.redbooks.ibm.com/redbooks/pdfs/sg245352.pdf

[38] D. Box, D. Ehnebuske, G. Kakivaya, A. Layman, N. Mendelsohn, H. F. Nielsen, S. Thatte, and D. Winer, "Simple Object Access Protocol (SOAP) 1.1," 5 2000, w3C Note 08 May 2000. [Online]. Available: http://www.w3.org/TR/2000/NOTE-SOAP-20000508/

[39] Twitter, "The Twitter REST API," 2014. [Online]. Available: https://dev.twitter.com/docs/api

[40] P. Devijver and J. Kittler, *Pattern Recognition: A Statistical Approach*. Prentice/Hall International, 1982.

[41] G. J. McLachlan, K.-A. Do, and C. Ambroise, *Analyzing Microarray Gene Expression Data*. Wiley-Interscience, 2004.

[42] S. Muller, "ODP makes moot the question, 'Do you support that protocol?'," 10 2014. [Online]. Available: https://communities.ca.com/community/ca-devtest-community/blog/2014/10/13/odp-makes-moot-the-question-do-you-support-that-protocol

[43] D. Swan, "Swinburne, CA Technologies in virtualisation breakthrough," *The Australian*, December 16 2014.

[Online]. Available: http://www.theaustralian.com.au/business/latest/swinburne-ca-technologies-in-virtualisation-breakthrough/story-e6frg90f-1227158041554

[44] V. J. Hodge and J. Austin, "A survey of outlier detection methodologies," *Artificial Intelligence Review*, vol. 22, no. 2, pp. 85–126, 2004.